

REMARKS

Claims 1-38 are pending in this Application of which Claim 1 is the independent claim. All claims stand rejected.

Claim 1 is being amended to further clarify the Applicant's invention, by reciting that "each JPU in the second group" is "responsive to requests received from a host computer to execute jobs and return results independent of execution at other JPUs." This amendment clarifies the preceding limitation that the "Job Processing Units (JPUs)" operate "autonomously and asynchronously from one another." Support for this amendment is found at least on page 11, lines 7-12 and 22-27 of the Specification as originally filed, where it is stated that "individual JPUs are free to complete the sequence on their own data without coordinating with other JPUs, waiting for results from other JPUs, or otherwise being constrained in the timing of their completion." Acceptance is respectfully requested.

Applicants note with appreciation the reconsideration and withdrawal of the objections to the Specification, as well as the double patenting rejections, as indicated in the instant Office Action.

Rejection of Claims 1-38 under 35 U.S.C. § 102(e)

Claims 1-38 have been rejected under 35 U.S.C. § 103(a) as being unpatentable over Kabra et al. (U.S. Patent No. 6,507,834). Applicants respectfully disagree with this rejection and request reconsideration.

Claim 1 is directed to an asymmetric data processing system comprising a first group of nodes comprising one or more host processors, a second group of nodes comprising one or more Job Processing Units (JPUs), and a network connecting the nodes within each group. To aid in understanding the present invention, Applicants refer to an exemplary embodiment represented in Figs. 1 and 7 of their Specification. Here, a host computer 12 communicates with Job Processing Units (JPUs) 22-1 – 22-i via a network 34, where each JPU 22 may access data at respective storage devices 23.

In an example operation, the host computer 12 receives requests from a requester 20, 33, 36 (e.g., a client computer or application) to process data stored at a plurality of storage devices

23 (e.g. hard disk drives). A plan generator (204 at Fig. 3) at the host computer 12 generates a plan for processing the request (Specification, page 46, lines 21-27). The plan comprises a number of jobs, which are distributed among the JPUs 22-1 – 22-3 (Fig. 7) and host computer 12 (Specification, page 48, lines 24-27). Each job further comprises a sequence of instructions that are executed by the JPUs 22-1 – 22-3.

In completing these jobs, each JPU 22-1 – 22-3 operates autonomously and asynchronously from one another. (Specification, page 9, lines 1-14). For example, each JPU 22 may be configured such that it can complete requests without waiting from input from other JPUs. As a result, each JPU may schedule jobs without regard for scheduling at other JPUs, thereby enabling each JPU to optimize processing of requests (Specification, page 11, lines 7-12 and lines 22-27).

In the instant Office Action, section “Response to Arguments” (page 4), it is stated that Kabra discloses “a second group of two or more Job Processing Units (JPUs), operating autonomously and asynchronously from one another.” In support of this assertion, Examiner cites Kabra at col. 7, lines 19-26 and Fig. 1: “communication between processors on a symmetric multiprocessing system, memory used as the transport vehicle.” However, neither this citation, nor any other section of Kabra, teaches JPUs that are autonomous and asynchronous from one another. Claim 1 has been amended to clarify the limitation that the JPUs operate “autonomously and asynchronously from one another,” by further providing that “each JPU in the second group” is “responsive to requests received from a host computer to execute jobs and return results independent of execution at other JPUs.” This feature is described in further detail on page 11, lines 7-12 and 22-27 of the Specification as originally filed: “For example, if a group of JPUs are collectively assigned a sequence of jobs by the host, individual JPUs are free to complete the sequence on their own data without coordination with other JPUs, waiting for results from other JPUs, or otherwise being constrained in the timing of their completion. This frees individual JPU to then run other jobs that may even relate to other queries, while neighboring JPU’s continue to process jobs from the first query.” Further, “a JPU processes requests asynchronously and autonomously. Each JPU is thus free to process its requests as quickly as it can, and return its results (partial or complete) to the requester.” Thus, each JPU operates without being constrained by the timing or execution of another JPU.

In contrast, Kabra discloses data servers 130 that operate in parallel, and fails to disclose operating autonomously and asynchronously. As shown in Figs. 2 and 3, Kabra uses a parallelizer 202 to create a plan for parallel execution of a data query. The plan is distributed as “segments” among data servers 130 (col. 9, lines 27-30). These segments are “executed concurrently by different data servers 130,” and so “the parallelism of database system 100 is established” (col. 9, lines 16-20). Kabra only describes executing segments concurrently at the data servers 130 as above, and fails to teach or suggest the data servers 130 operating asynchronously from one another. For at least the above reasons, Kabra fails to disclose an “asymmetric data processing system” as recited in Claim 1 of the present Application.

Further, the Office Action cites Kabra at col. 7, lines 19-26 as disclosing “operating autonomously and asynchronously from one another” (Office Action, page 4, paragraph 3). However, Kabra at col. 7, lines 19-26 merely describes a network transport protocol (TCP), and the use of memory to transport data on a symmetric multiprocessing system (SMP). An SMP system relates to multiple processors sharing a memory, which typically requires that the processors be synchronized (not asynchronous) in order to communicate with the memory. Thus, an SMP system does not provide for autonomous or asynchronous operation. Rather, Kabra’s data servers 130 operate concurrently and in parallel, as described above. For at least the above reasons, Kabra fails to disclose an “asymmetric data processing system” having “two or more Job Processing Units (JPUs), operating autonomously and asynchronously,” as recited in Claim 1 of the present Application.

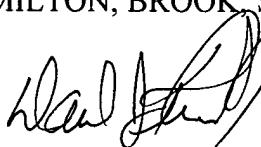
Claims 2-38 each depend from Claim 1 and thus are allowable at least for the reasons stated above. As a result, the §102 rejection of Claims 1-38 is believed to be overcome, and reconsideration is respectfully requested.

CONCLUSION

In view of the above amendments and remarks, it is believed that all claims are in condition for allowance, and it is respectfully requested that the application be passed to issue. If the Examiner feels that a telephone conference would expedite prosecution of this case, the Examiner is invited to call the undersigned.

Respectfully submitted,

HAMILTON, BROOK, SMITH & REYNOLDS, P.C.

By 
David J. Thibodeau, Jr.
Registration No. 31,671
Telephone: (978) 341-0036
Facsimile: (978) 341-0136

Concord, MA 01742-9133

Date:

5/22/08